Parallel Algorithms HW 4

Sparse matrix-vector multiplication

1. In C implement a sparse matrix-vector multiplication algorithm giving $b = A.x$. b and x are stored as standard single precision vectors and (presumably sparse) A is stored by row with each row defined by two lists: A list of integers with the location of each non-zero and a list of floats (single precision) giving the values of those no-zeros.
2. Test you C code on a variety of small problems.
3. In CUDA implement a sparse matrix-vector muliplication algorithm with each thread responsible for computing one entry of the output vector b using the same storage scheme.
4. Test you CUDA code on a variety of small problems.